Providing Location-Independent Access to Patient Clinical Narratives Using Web Browsers and a Tiered Server Approach

J. Craig Klimczak, D.V.M., M.S.[1,2], David M. Witten, II[2], Michael Ruiz, R.N., M.B.A.[3], Joyce A. Mitchell, Ph.D.[2], John G. Brilhart, M.B.A., C.D.P.[3], and Michael L. Frankenberger[3,4]
[1]Program in Health Services Management, [2]Medical Informatics Group, School of Medicine, [3]Information Services University of Missouri Hospital and Clinics, [4]Chief Information Officer, University of Missouri - Columbia, MO 65211
(573) 882-2718

ABSTRACT

Health care today depends upon timely access to patient medical data and the latest medical knowledge. As we make the transition from a hospital-based organization to an integrated health care delivery system, patient care information must move throughout the organization quickly and efficiently over increasing distances. The emergence of widely-dispersed referral networks demands novel solutions to the problems of delivering patient care information to providers. We have developed a mechanism to provide location-independent access to clinical narrative reports using a multi-tiered server model and World Wide Web technologies for delivery. To successfully deploy such a system to sites separated by large distances, it is important to reduce complexity at the client site. Using a "thin client", such as a web browser, in our design facilitates deployment and support while reducing cost per user. This architecture allows the application to be updated without modification to the end-user software and eases maintenance over long distances.

INTRODUCTION

Despite advances in computer technology, the medical record continues to be paper-based in many health care institutions, including our own. In a recent study at a university hospital clinic, Tang, et al. found that in 81 percent of the cases, physicians could not find all of the patient information needed during a patient's visit.[1] They concluded that the multi-authored paper medical record did not provide effective access to patient information for physicians making clinical decisions.

In another study, Hripcsak reported "Much clinical data is locked up in departmental word-processor files, clinical databases, and research databases in the form of narrative reports such as discharge summaries, radiology reports, pathology reports, admission histories, and reports of physical examinations. Untold volumes of data are deleted every day after word-processor files are printed for the paper chart and for mailing reports."[2] The situation that Hripcsak described at Columbia-Presbyterian Hospital exists at the University of Missouri Hospital and Clinics (UMHC) as well. While this use of information technology has improved the efficiency in which dictated documents are transcribed, printed, and placed in the paper chart, it has done nothing for the evolution of the patient medical record nor for solving the problems that we encounter everyday with the paper-based patient chart.

Health care today depends upon timely access to patient medical data and the latest medical knowledge. The process of assembling, producing, and delivering a health care service to a patient involves the intensive use of information and the "product" contains high information content. Information is, in fact, the essential raw material to be processed and transformed in the delivery of health care. Much of this information is supplied by the patient, acted upon by the health care provider, and stored away in the patient's chart for future use within the organization. Health care providers need location-independent, real-time, secure access to this information.

Not only is this information needed by health care providers, but it is also needed by office support personnel. Hospital operations and management are as dependent upon access to patient information as any physician or health care provider. To a hospital, the patient chart is not a set of notes maintained by a physician for patient care, but rather a comprehensive log of all patient interactions occurring during that patient's visit to the hospital. This log plays not only a vital role in patient care but also in assuring that a hospital meets mandated documentation guidelines for reimbursement and accreditation. Patient chart information is used by many departments throughout the hospital to establish fees, generate invoices and ensure quality service.

Unfortunately, no mechanisms exist to provide access to these documents other than the printed report and chart. Further, there is no index to these documents that one could use to find information on a particular patient. The failure of the medical record to evolve away from the paper chart leaves us a legacy of problems that health care providers must face every day.

PROBLEMS WITH THE PAPER CHART

There are at least 4 problems associated with the paper-based medical record: 1) access, 2) vulnerability, 3) overhead, and 4) timeliness.[3] Each of the 4 problems is distinct and offers special considerations for a text archive and retrieval system.

Access to the paper chart is limited to the physical location of that chart. When multiple groups need access to the chart, numerous copies must be printed and distributed. This results in wasted paper, delays, and excessive utilization of manpower to copy, address, route, and transport these paper documents to their respective locations. When these charts get lost or misplaced, personnel must manually search for the missing information. Turnaround time from dictation to placement in the chart can be several days despite the fact that most transcription is completed in less than six hours. Frequent updates and revisions lead to out-of-date information which creates potentially harmful situations for patients.

ELECTRONIC HEALTH CARE RECORD

Many believe that the electronic health care record (EHCR) will provide the solution to these problems. Since much of the clinical information currently exists in legacy information systems, we must develop technologies that integrate information contained in these systems with new information systems and provide a means to securely access it.

IAIMS at the University of Missouri

The inception of the EHCR project began with an Integrated Advanced Information Management System (IAIMS) planning grant from the National Library of Medicine. This grant fostered the formation of the Missouri Integrated Advanced Information Management System (MIAIMS) project. Under the MIAIMS project a Patient Care Planning Subcommittee was formed to:

1) Analyze the information needs for patient care management;

2) Identify factors affecting future support and continued expansion of information resources to achieve the information needs; and

3) Prepare a report for information system development, integration, and implementation with supporting documentation.

This committee established the development of an electronic health care record as one of the institution's top priorities. The committee analyzed the patient chart and found that much of the content was transcribed clinical narratives. Since laboratory results, pharmacy records, orders, and radiology reports were already online in departmental systems and clinical narratives were not, the committee chose to initiate the System for Text Archive and Retrieval (STAR) project as the first step toward an EHCR.

REQUIREMENTS

The STAR project began with the development of a prototype system by the Medical Informatics Group. The prototype was constructed in Visual Basic™ using a client-server model. From this prototype a set of system requirements was written. These requirements became the basis for a request for proposals (RFP) that was used to select a vendor partner. Infodata Corporation of Fairfax, VA was selected to assist refinement of the requirements, development of a pilot system, and implementation of STAR.

Many goals and requirements were defined for the STAR project. These included:

1) the ability to edit and complete clinical narratives online;

2) the ability to electronically sign and authenticate clinical documents online;

3) the ability to route, track, and queue documents for office processing and work flow management;

4) the ability to retrieve patient reports based on structured data found in current patient management systems (i.e. by attending physician, by primary care physician, by referring physician, by patient name, number, etc.);

5) the ability to retrieve patient reports based on information contained within the reports themselves either by word or concept indexes;

6) improvement of the efficiency of communications between affiliated referring physicians and in house providers;

7) use of common off-the-shelf (COTS) software where available;

8) provision of tight integration with legacy and library systems for one stop access to information;

9) provision of a security mechanism that supports a policy that a provider/patient relationship must

exist for a particular provider to be allowed to view a given patient's document.

10) provision of a security mechanism that would increase the anonymity of retrospective chart analysis minimize the impact on patient privacy; and

11) use of a scaleable architecture that could support up to 2000 users.

Making the transition from a prototype to a fully functioning system illustrated that many of the goals and requirements of the project could not be achieved on a first pass. Several of the planned and previously demonstrated features would have to be postponed until later versions. Working with the team from Infodata the UMHC design group decided to scale back the requirements for the pilot project and concentrate on architectural issues.

SYSTEM ARCHITECTURE

Two architectural dimensions became apparent to the design team during the process: 1) where the documents would be stored and 2) whether the integration of data would occur on the client or at the server. This latter consideration dictated whether there would be a single interface to the server or whether there would be multiple interfaces to separate sources of data.

The goals of the prototype were to prove the concept of the Text Archive System and to test the client server design. In an effort to have a minimal impact on the feeder systems a leave-in-place design was adopted (Figure 1).

An Oracle™ database server, a Fulcrum™ text search server, and a Novell™ file server were incorporated into this application,. Using 3 servers required separate connections from the client to each of them. When integration of data was required, it was performed at the client. For example, the client would request a list of patient numbers for a particular doctor from the Oracle™ database and would use that list of numbers to retrieve documents from the text server. All program logic was contained at the client. This model of client is sometimes referred to as a "fat client".

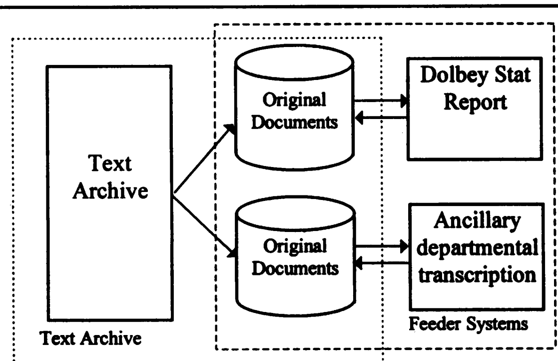

Figure 1: Leave-in-place design

During the course of development many problems were identified with the "fat client" model. We found that the client software was extremely difficult to install on new workstations. Installations were complicated by the number of drivers, network support layers, and link libraries that were needed. A slight change in configuration to any of these components would render the system inoperable.

These problems led us to look at alternative implementation strategies. First, we switched to a design that copied the documents to a separate archive (Figure 2). This allowed us to design a single interface to the text archive for the feeder systems. This eliminated the dependence on the feeder's file system for document storage and provided greater control over the collection.

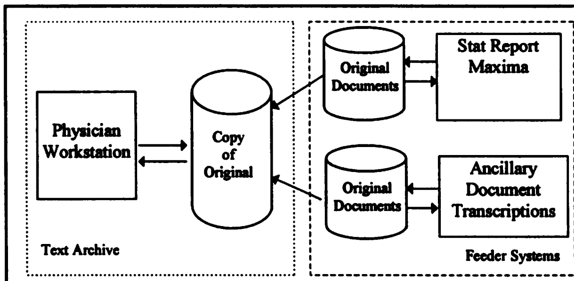

Figure 2: Copy to separate archive design

Second, we switched to a design where the intelligence of the system was built on the server side of the application rather than on the client side. In this design, we employed a Web browser as a "thin client". This approach permits the application to be "pushed" from the server to the client. The use of this design pays dividends during system maintenance and upgrades because changes need only be made on the server. This eliminates the need to update potentially thousands of copies of the client application residing on geographically-distributed client workstations.

Next we had to deal with the issues of integrating the data from multiple sources. This included patient demographic data from 3 disparate registration systems and document profile data from the feeder systems. This information had to be made searchable from the Web client to identify specific documents.

Interfaces were installed that would allow real-time transfer of this data to the archive's Oracle™ database. In addition, we designed a data model that would permit merging of the disparate registration systems' information in the archive's database. Techniques were developed to link the demographic data from the registration systems to document profiles and the documents themselves.

Using Web browsers as clients requires a Web server at the host site. Since transaction-orientated database services are not a standard part of Web servers, a multi-tiered server design was needed to integrate the relational database and document archive with the Web Server. A custom-designed interface layer, which uses the Netscape™ API and the Common Gateway Interface (CGI) to communicate with the Web server, was written to control access to the databases and document collection. This intermediate tier allowed result sets from a query to the relational database to be passed to the free text search engine. A single combined result is sent to the client. Tiered servers allow the use of a wide variety of interface technologies to combine information from multiple sources.

The multi-tiered server model provides many advantages to Web-enabled applications. First, it provides a mechanism to maintain the state of a user's connection. Web servers by themselves are stateless and retain no information about a users' request beyond returning the page requested. There is no continuity between individual requests. On the other hand, database servers are session-orientated. Second, the multi-tiered server model provides an essential layer of security that is needed in a health care environment. In our system, each document request is an indirect query to the intermediate interface layer which authenticates access to the collection. This prohibits unauthorized users from directly retrieving patient reports from the collection. Since documents are not directly addressed by the Web server, bookmarks and spoofing attacks cannot be used to obtain access to the collection. An additional layer of privacy is provided by the Netscape Commerce Server secure sockets layer which encrypts potentially sensitive patient information as it is being transported across public wide area networks. This is a critical design feature for the STAR project since we are using the system in referring physician practices that are attached to public Internet connections.

Figure 3 illustrates the multi-tiered server design. At the top of the figure is the Netscape Commerce Server that is used to manage requests from the user's Web browser. Between the Web Server and databases is the intermediate interface layer that processes CGI requests and passes them to the database server layer. The database servers process the requests and results are passed back to the intermediate layer, which then assembles the results and returns them to the Web

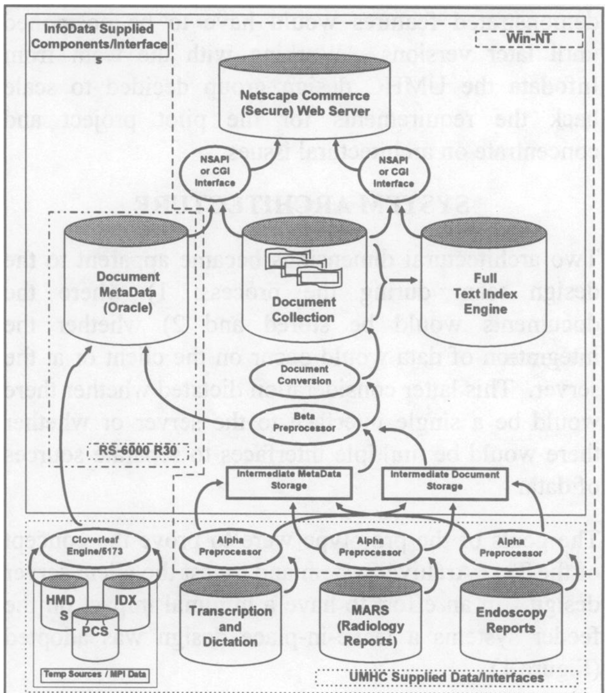

Figure 3: Multi-tiered server design of STAR project

server.

The database server layer is composed of three separate databases providing separate services. A relational database server is used to manage the fielded data provided by the registration systems and document profile data from the feeder systems. A free text index server provides free text searching capabilities. The archive's operating system is currently used to store and retrieve the document files. Storage of the documents in a relational database has been considered but not implemented.

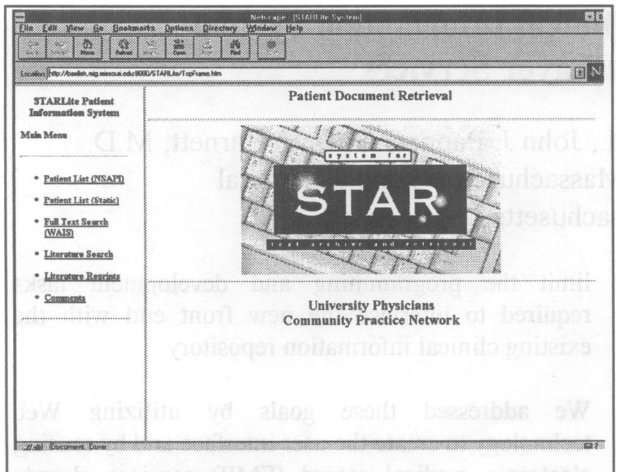

Figure 4: Opening screen and menu

USER INTERFACE

Physicians access documents by navigating through a series of displays presented by the Web browser. Any Web browser that supports Hypertext Markup Language (HTML) version 2.0 or greater can be used. Once a physician is verified as an authorized user, the first of these displays, the Patient Search Form, is presented. The physician can fill in any combination of fields on the search form and submit a search request to the server. The server responds by sending the physician a list of patient names to choose from. A physician can either request additional patient information, which is displayed in a different frame on the right side of the screen, or request a list of documents for a patient, which are displayed on a different screen along with patient identification information.

The document list display is broken up into patient identification information, document viewing, and

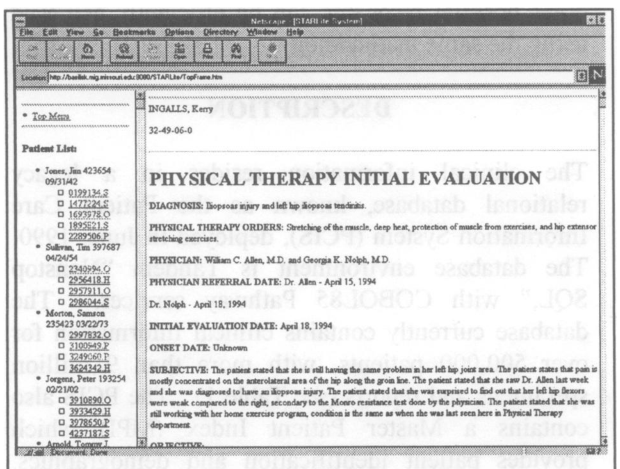

Figure 5: Example physical therapy note

document list frames. To view an individual document, the physician simply chooses one of the documents from the document list frame. The document is retrieved from the server and displayed in the frame on the right (Figure 5).

SUMMARY

Paper medical records are fraught with problems relating to access, vulnerability, overhead, and timeliness. The potential exists to solve many of these problems with an Electronic Health Care Record. The need for an EHCR is exacerbated by the changes that are occurring in health care today. As managed care pushes health care providers toward integrated delivery networks, information systems will be needed to provide communications and continuity across the organization. Health care providers will need secure, location-independent, real-time access to the patient's medical record.

Providing this type of access over large geographic regions is a technological challenge. Many approaches and models have been used to meet this challenge. However, the client-server model seems to be the most attractive. Client-server models are not without their problems. While "fat clients" have demonstrated superior performance and features, they are often handicapped by fragile configurations and complicated installation procedures. The World Wide Web paradigm coupled with a multi-tiered server approach offers a "thin client" solution that is scaleable and simpler to maintain.

Acknowledgments

This work is supported in part by grants LM05415 and LM07089 from the National Library of Medicine.

References

[1] Tang PC, Fafchamps D, Shortliffe EH. Traditional Medical Records as a source of Clinical Data in the Outpatient Setting, in Ozbolt JG, Proc. 18 SCAMC: Washington D.C. 1994;18:575-579.

[2] Hripcsak G, et. al.: Unlocking Clinical Data from Narrative Reports: A study of Natural Language Processing. Ann Intern Med. 1995;122:681-688)

[3] Klimczak JC, et. al.: Reengineering Medical Records With a Text Archive and Retrieval System. In Proc. 1996 Annual HIMSS Conference: Atlanta, GA 1996;3:64-76.